12.7 Cascading Techniques for a High-Speed Memory Interface

Zheng Gu, Peter Gregorius, Daniel Kehrer, Lydia Neumann, Evelyn Neuscheler, Thomas Rickes, Hermann Ruckerbauer, Ralf Schledz, Martin Streibl, Juergen Zielbauer

Qimonda, Munich, Germany

As the data rate exceeds 3Gb/s, the usability of the present single-ended memory interface with multi-drop configuration is reaching its limit due to increasing signal-integrity issues. Differential signaling combined with point-to-point connection is known to reduce cross-talk and signal reflection [1], but a point-to-point connection limits the number of DRAM components in the system. A repeater function is a suitable approach to enable higher memory densities [2]. A typical example of such a memory interface is the advanced memory buffer used in fully buffered DIMM for server applications [3].

Similar signaling concepts may be integrated into future standard DRAM components. A major challenge for a high-speed memory interface will then be a cost- and power-efficient repeater function. There are two basic concepts, a transparent repeater mode without retiming of the transferred data and a resample mode with retiming of the transferred data. The resample mode requires a continuously running low-jitter clock and hence consumes significant power. In contrast, the transparent mode repeater allows the clock to stand by, which significantly reduces power consumption. However, there are substantial jitter-accumulation and signal-integrity concerns in such a system.

This has motivated a test chip to assess the concept of a transparent-mode repeater. The chip consists of 6 parallel repeat lanes for data and a single repeat lane for clock (Fig. 12.7.1). The clock trunk extends over 6.5mm to account for realistic distances expected for a DRAM pad row. The packaged chips are placed on a PCB with a spacing of 10mm. The chip is programmable by a serial interface with 64 control bits per lane for performance tuning and test.

A clock-generation block (Fig. 12.7.1) with a CML-type frequency divider and phase interpolator receives the half-rate clock from the clock trunk and provides 4 quarter-rate CMOS clock phases with a control resolution of 6 degrees (1/15 UI). The differential receiver (Fig. 12.7.2) allows for offset compensation and performs receive equalization by gain peaking that is adjustable via a MOS capacitor array. Data is sampled by 4 interleaved sense-amplifier/SR-latch type sampling stages including a ±30mV offset compensation. The 4:1 serializer consists of a pulse generator for regeneration of the 1UI time base [4] and a single-stage 4:1 multiplexer circuit that improves clock phase margin and power versus a tree-like serializer. The CML-type selector forwards 1 out of 3 signal sources: the transparent repeater path, the resample path, or a ring oscillator used as performance monitor. The CML type pre-driver with active loads has a replica-bias circuit (Fig. 12.7.3) and an output swing control. The TX driver (Fig. 12.7.3) is a CML buffer with a regulated current source that allows lower width of the current-source transistor compared to a conventional current mirror. The tail current of pre-driver and TX driver can be adjusted via control bits. The ODT is adjustable from 40Ω to 80Ω.

The circuit is fabricated in a 3M 70nm standard DRAM process [5]. A DRAM-typical single-layer wire bond BGA package is used. The chip operates at 1.35V supply and occupies 6.9×0.5mm^2. Figure 12.7.7 shows a micrograph of one of the data repeat lanes. The single lane occupies 0.45×0.36mm^2.

A BER tester is used as data generator and analyzer. Sampling and real time oscilloscopes with 12GHz differential probes are used to evaluate data eye quality in the repeater chain. The device showed robust functionality up to 6.4Gb/s in all operation modes on a test board for single component characterization. The system test board with a repeater chain of 6 ranks is implemented on standard low-cost FR4.

To assess the maximum applicable number of transparent repeats, input eyes and jitter accumulation have been probed at inputs of ranks 1, 3, and 5 (Fig. 12.7.4). The eye diagrams show double traces due to bandwidth limitations of channel and package. Resulting bathtub traces for ranks 1, 3, and 5 are given in Fig. 12.7.5. Measurements at 4.8Gb/s show an eye degradation of ~8% UI per rank due to jitter accumulation.

To prove that transparent repeat data is correctly recovered by an arbitrary rank, all but one of the ranks 0 to 5 are set to transparent repeat mode, while the one rank is used to sample and regenerate the data. The effective eye opening captured by the samplers is smaller than the eye opening probed at the RX (Fig. 12.7.5) due to clock jitter, clock phase mismatch, and setup/hold time of the sampler. In addition, power-supply noise is added by power-up/power-down cycles in neighboring lanes with a frequency of 200MHz. Measurements still show a comfortable sampler eye opening of 0.25UI for a BER of 10e-12 at Rank 5 (Fig. 12.7.5). Close to the repeater performance limit at 5.3Gb/s, the RX gain peaking helped significantly to improve the timing margin for the sampler and gave an error-free operation at all ranks during 2e14 transmitted bits at 5.3Gb/s for a PRBS-32 sequence.

In contrast to the transparent repeat path, the resample mode is affected by clock phase mismatch (Fig. 12.7.6). Monte-Carlo simulations predict a 7% standard deviation of the eye opening due to mismatch of the quarter-rate clock phases. Such timing distortions can in fact be found on silicon (Fig. 12.7.6). The chip consumes 50mW per lane for the transparent mode compared to 80mW for the resample mode. Moreover, the resample mode increases latency by 4UI functional and 390ps propagation delay per rank.

A transparent repeater chain of 6 ranks, based on a standard DRAM process and a standard wire-bond package is demonstrated to achieve up to 5.3Gb/s with a BER<1e-14. At 4.8Gb/s, the last rank in the chain has an excellent RX eye opening of 0.5UI at BER<1e-12. The transparent repeater mode consumes 40% less power and has 80% less latency compared to the resample mode. This provides an effective option to increase memory density for a differential point-to-point high-speed interface.

Referencees:

[1] S. Lee, K. Kwon, C. Kim, et al., "A 128Mb Multi Port Media DRAM with four Independent 4Gbps Serial Ports," *Symp. VLSI Circuits Dig.*, pp. 34-35, Jun., 2004.

[2] J. Kennedy, R. Moorey, R. Ellis, et al., "A 3.6-Gb/s Point-to-Point Heterogeneous-Voltage-Capable DRAM Interface for Capacity-Scalable Memory Subsystems," *IEEE J. Solid-State Circuits*, vol. 40, pp. 233–244, Jan., 2005.

[3] H. Partovi, W. Walthes, L. Ravezzi, et al., "Data Recovery and Retiming for the Fully Buffered DIMM 4.8Gbps Serial Links," *ISSCC Dig. Tech. Papers*, pp. 336-337, Feb., 2006.

[4] D. Kehrer, H. D. Wohlmuth, "A 30-Gbps 70-mW One-Stage 4:1 Multiplexer in 0.13µm CMOS," *IEEE J. Solid-State Circuits*, vol. 39, pp. 1140–1147, Jul., 2004.

[5] J. Amon, A. Kieslich, L. Heineck, et al., "A Highly Manufacturable Deep Trench Based DRAM Cell Layout with a Planar Array Device in a 70nm Technology," *IEDM Dig.*, pp. 73–76, Dec., 2004.

 1-4244-0852-0/07/$25.00 ©2007 IEEE.

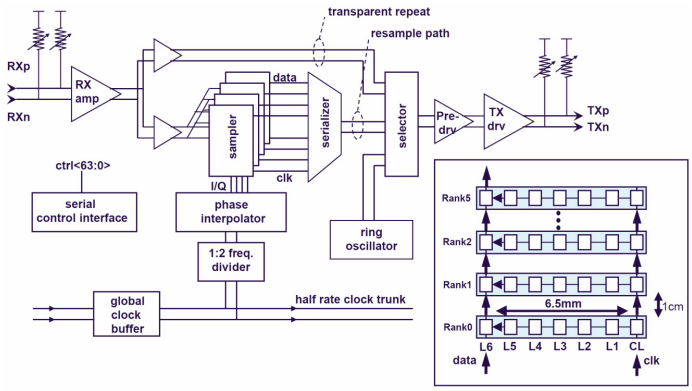

Figure 12.7.1: Block diagram of a single high-speed lane. The inset shows the test setup with 6 chips, each chip having 6 data repeat lanes and a clock lane.

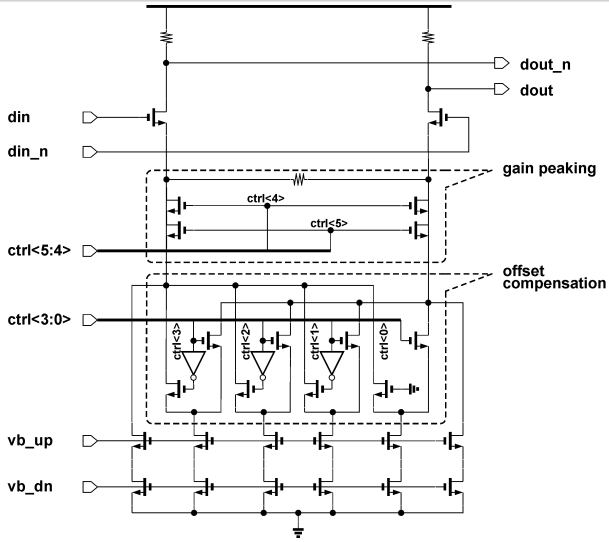

Figure 12.7.2: Schematic of RX-amplifier with bandwidth enhancement and offset compensation.

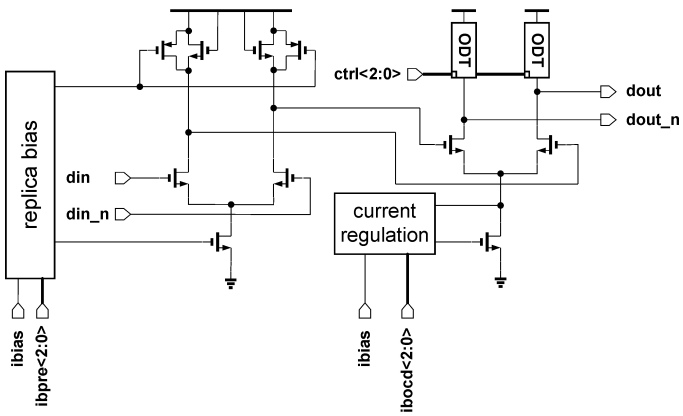

Figure 12.7.3: Schematics of pre-driver and TX driver.

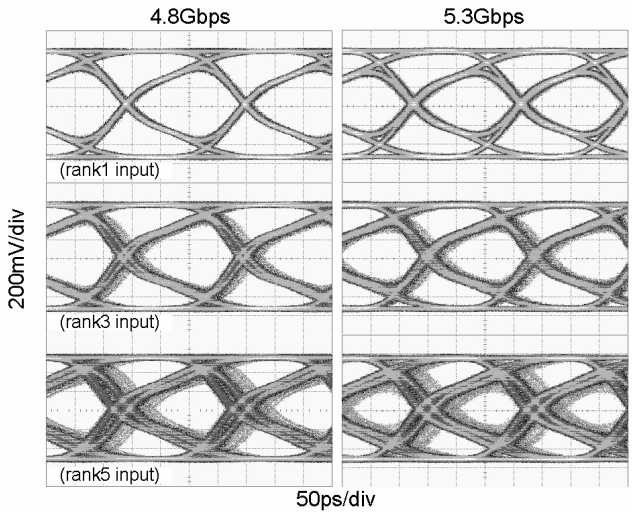

Figure 12.7.4: Measured eyes at ranks 1, 3, 5 at 4.8Gb/s and 5.3Gb/s (all ranks in transparent repeat mode).

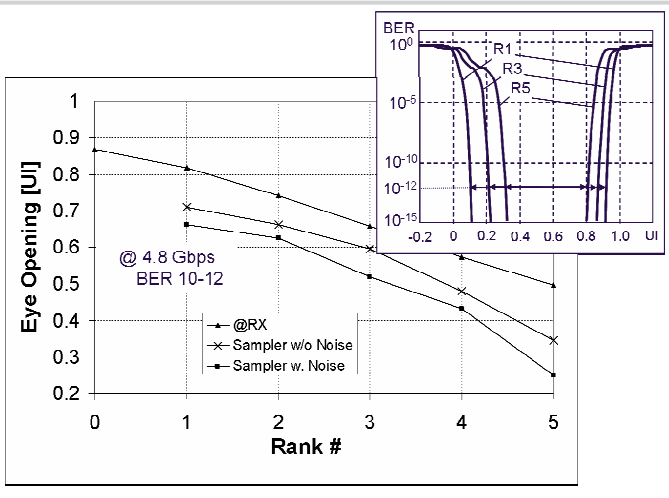

Figure 12.7.5: Eye opening captured by the sampler with and w/o noise. Eye opening probed at input of ranks 0 to 5 and corresponding bathtub traces for the RX eye of ranks 1, 3, and 5.

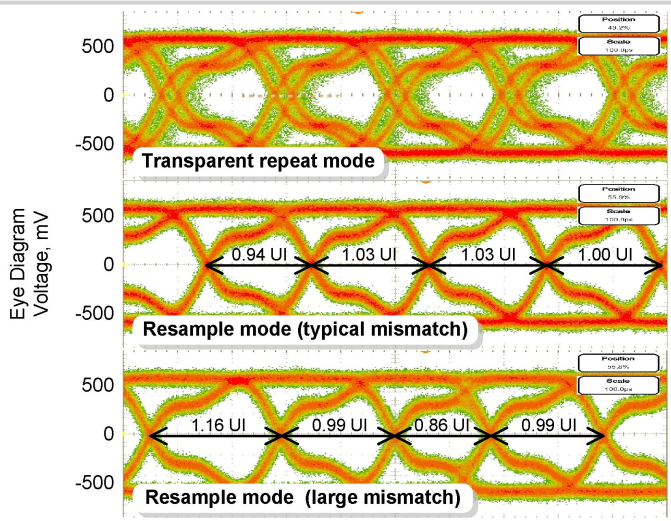

Figure 12.7.6: Measured eye diagrams at input of rank 5 @4.8Gb/s. Top: transparent repeat mode; middle and bottom: different chipsin resample mode (device mismatch).

Continued on Page 599

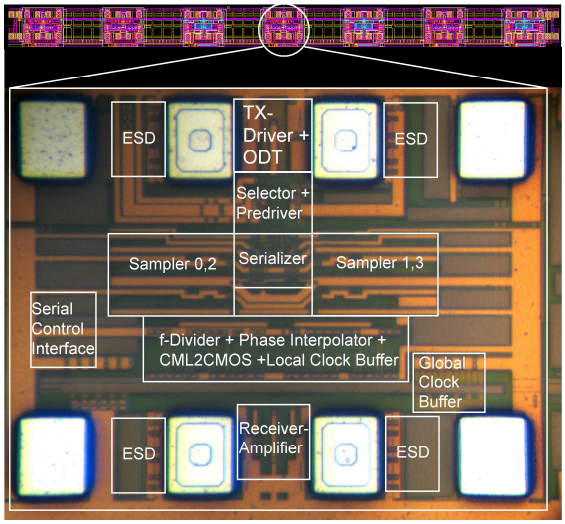

Figure 12.7.7: Layout of the 6.9×0.5mm² chip (top) and chip micrograph of a 0.45×0.36mm² high-speed lane (bottom).

 1-4244-0852-0/07/$25.00 ©2007 IEEE.